Future Proofing the DMA for Agentic AI: Lessons from the AI Act

Jan-Frederick Göhsl*

*Assistant Professor for Civil Law, Economic Law and Law and Digitalization.Email: jgoehsl@muenster.de

Article Info

How to cite this paper: Jan-Frederick Göhsl. (2025). Future Proofing the DMA for Agentic AI: Lessons from the AI Act. *World Copmpetition, 48*(1), 10-50.
https://doi.org/10.54648/woco234290

ISSN Online: 1875-8436
ISSN Print: 1011-4548

Received: 27.02.2024
Revised: 12.06.2025
Accepted: 01.03.2025
Article Type: **ORIGINAL Research**

ABSTRACT

The Digital Markets Act (DMA) was adopted to address known (competition) issues in digital markets. Its effectiveness with regard to emerging technologies – i.e., its future-proofness – is therefore an open question. In this article, we tackle that question via a case study of AI agents. AI agents are the most consequential emerging technology on the horizon, with the potential both to disrupt existing gatekeepers and to create gatekeepers in their own right, but policymakers did not consider the inclusion of AI (agents) when designing the DMA. We start by explaining the key building blocks of agentic AI and by identifying potential competition concerns. Next, we explore whether AI agents could fall within the DMA's scope (as core platform services) and whether the ensuing obligations are meaningful when applied in this new context. We find that the DMA is surprisingly future-proof in the face of agentic AI, but also make recommendations on (re)interpretations and amendments of the DMA to ensure its future effectiveness and legal certainty.

Keywords: Artificial Intelligence, AI Agents, Agentic AI, AI Act, Digital Regulation, Digital Markets Act, Core Platform Services, Gatekeeper, General Definition, Delegated Acts.

1.Introduction

The Digital Markets Act (DMA) is the primary EU instrument that regulates competition in digital markets. It was adopted in 2022, in large part to make up for the limitations of competition law. Instead of competition law's ex post enforcement, the DMA institutes an ex ante regime whereby gatekeepers in a number of 'core platform services' (CPSs) are first designated as such and then subjected to a list of obligations ('dos and don'ts'). The European Commission ('Commission') designated gatekeepers in 2023–24, after which gatekeepers started implementing the obligations. By 2025, the Commission had adopted its first non-compliance decisions.4

It is an open question how future-proof the DMA is. On the one hand, the DMA is largely designed to solve issues identified in previous competition law enforcement – in other words, problems of the past. In this sense, its future-proofness could be a secondary consideration. On the other hand, the DMA explicitly expresses lawmakers' desire to "ensure that this Regulation

One challenge is that AI is difficult to fit within the DMA's logic. AI is a general-purpose technology, rather than a specific platform that serves as a gateway between business and endusers, and that would thus fit on the list of CPSs. In other words, the DMA is about applications more than it is about technologies, and AI is more the latter than the former. One specific application of AI is currently found in chatbots, but their competitive significance appears limited. Recently, however, we have observed the rise of an AI-based service that does have the potential to transform markets: AI agents.

Unlike chatbots/AI assistants (such as ChatGPT and Copilot), which primarily respond to user inputs, AI agents exhibit proactive behaviour. Following an initial prompt, they can continue their work without further input or human intervention, and may even plan ahead and execute actions autonomously. Agentic AI systems possess persistent memory, enabling them to adapt and learn from their experiences. They possess enhanced connectivity, allowing them to accept (multimodal) input from various sources (e.g., voice, text, vision, sensors, or external applications) to effectively comprehend and carry out tasks. Agentic AI systems seamlessly integrate with external applications and tools, facilitating even communication, invocation, and collaboration with other AI agents.

While this may sound like science fiction to some, the first prototypes of agentic AI systems have already been showcased by players like OpenAI (first called "Operator", now "ChatGPT Agent"), Google (DeepMind's "Project Astra"), Anthropic (Claude's "Computer Use"), and Perplexity (its browser-based "Comet"). Currently, most prototypes of AI agents are based on an "act-like-a-human" architecture, where the AI agent can "see" what the user typically sees (by capturing the user's interface, e.g., the browser window or the entire user's screen) and can perform actions through clicking or typing just like a user would. By mimicking human

interaction, the AI can make use of existing interfaces designed for humans. In turn, this does not require application programming interfaces (APIs) specifically designed for AI agents, or

additional permissions and entitlements. However, this is likely only a first step in the evolution of AI agents. In the future, AI agents will become more deeply embedded into technology, and technology and interfaces are likely to become more agent-centric. As AI agents advance, and their use becomes more widespread, new paradigms and forms of interaction will emerge. The agentic revolution has the potential to redefine industries, disrupt traditional business models, and create a new era of human–machine collaboration.

In particular, agentic AI may radically transform how consumers interact with their (mobile) devices, potentially leading to new interface designs and wearable devices, where the AI agent – rather than a specific device or app – provides the unique access point for users to discover and use digital services. This has the potential to disrupt the gatekeepers that are controlling today's access points like search, app distribution or operating systems (OSs). First signs of such new designs of mobile devices are already visible today. For example, in the vision of an "app-less" AI phone, the user gives tasks (e.g., to book a trip) to the AI agent, and it is the AI agent – rather than the user – that chooses which underlying apps or providers are employed to fulfil that task. The power to steer consumer demand, and even to act on behalf of the user, thus lies with the AI agent, which can therefore become a gatekeeper in its own right.

AI agents provide the perfect test case for the DMA's future-proofness. They likely constitute the most consequential platform or application that has emerged since the adoption of the DMA. AI agents have the potential to disrupt existing gatekeepers, which aligns with the DMA's goal of "contestability". Indeed, if contestability is to be found anywhere, it is likely in new market segments, rather than those where gatekeepers are already entrenched. Over time, agents could turn into gatekeepers themselves, at which point they can threaten the DMA's goal of "fairness". However, little policy attention has been paid to AI agents. Instead, competition authorities have focused on foundation models, as attested by reports by the CMA and Autorité de la Concurrence. Far less research has been devoted to the issue of agentic AI, which is situated further downstream and has unique competitive dynamics.

In this article, we therefore explore whether the DMA, which was not drafted with AI (agents) in mind, is nevertheless fit for the advent of this emerging technology. Specifically, we analyse whether agentic AI could be subsumed under any of the existing CPS categories, and whether the ensuing obligations are meaningful in this context. We conclude that the text of the DMA can be interpreted in such a way that it is surprisingly well-prepared for the new phenomenon of AI agents. Gatekeeper OSs, the likely chokepoint for AI agents, have already been designated, and certain CPS categories – in particular that of "virtual assistants" – allow for the designation of AI agents themselves, at least before they reach a high level of autonomy. The according DMA obligations can already address many of the potential competition concerns related to both the foreclosure of an emerging AI agent, in particular by an existing OS gatekeeper, as well as the more distant foreclosure of other services through a DMA-designated AI agent. However, we also make recommendations on (re)interpretations and amendments of the DMA that make it even more future-proof in the context of agentic AI.

The remainder of this paper is organised as follows. In Section II, we provide technical background on the definition and main building blocks of agentic AI. In Section III, we establish the main potential competition concerns that arise in the context of agentic AI, which fall into two categories: foreclosure of an (emerging) AI agent and foreclosure through a (gatekeeping) AI agent. In Section IV, we analyse whether and in which CPS category AI Agents could potentially be designated under the DMA. In Section V, we check how meaningful the DMA's obligations are when applied to AI agents and in light of the different designation possibilities. In Section VI, we evaluate whether the DMA is a "holistic regulatory response" to address concerns related to foreclosure of AI agents and foreclosure through AI agents. In Section VII, we conclude by making policy recommendations on how the DMA can be improved in the context of agentic AI.

2. BUILDING BLOCKS OF AGENTIC AI

Agentic AI has recently become a buzzword, but it is not yet well-defined. In this section, we offer our understanding of agentic AI, and its main building blocks, which we rely on for the remainder of the article.

Agentic AI is distinct from GenAI, although it employs GenAI. The primary defining characteristic of agentic AI is its greater degree of autonomy compared to other AI forms. It can and will make decisions independently, with minimal human oversight. This is also what differentiates AI agents from AI assistants, such as ChatGPT. AI assistants based on large language models (LLMs) are by now very familiar to us, but AI agents (will) far exceed their functionalities. However, as AI assistants currently undergo rapid development, we expect them to incorporate more and more agentic abilities in the near future, continuously blurring the line between AI assistants and AI agents (as seen in the introduction of an "Agent Mode" in ChatGPT). We believe that AI assistants will gradually transition into AI agents over time, which occurs on a spectrum defined by four key dimensions (i) the aforementioned degree of autonomy, (ii) the ability to adapt and learn from experience, (iii) perception of the environment, and (iv) the ability to plan and reason. We describe each of these dimensions in more detail below, while introducing the various "modules" that comprise the agentic AI system.

The degree of autonomy and human oversight involved will likely evolve gradually and be task specific. Currently, AI generally operates based on human-defined objectives. Fully autonomous AI systems that can set and pursue their own goals remain theoretical. While many AI systems exhibit some agentic behaviours, they are typically specialised for specific tasks and limited in scope due to safety and usability concerns. Agentic general AI capable of autonomous decision-making across a broader and more diverse range of tasks does not exist yet. The pursuit of such capabilities has sparked both interest and concern among researchers. Some experts raise significant safety concerns regarding potential risks, while others view agentic AI as the ultimate step towards artificial general intelligence (AGI). In essence, what distinguishes agentic AI from GenAI and other forms of AI (AI itself being a poorly defined and often loosely used term) is constantly evolving, contributing to the perception of the term "agentic AI" as fuzzy.

AI agents demonstrate not only greater autonomy but also have a greater ability to adapt and learn from experience. While so far GenAI has mainly been built on pre-trained models (that learned

from existing data), agentic AI learns more from experience and persistent memory. Some scholars have argued that data-driven network effects (positive data feedback loops) – such as associated with search engines, where users provide implicit feedback by clicking and viewing results – are not strong in the context of GenAI or AI assistants, precisely because this technology is using pre-trained models and does not learn from user input. Hence, a tipping of these markets – as in the case of search – is less likely in the context of markets for GenAI. However, this reasoning does not seem as relevant in the context of agentic AI, and indeed recent advancements in GenAI, which rely more heavily on more persistent memory and reinforcement learning, a machine learning technique where the model is improved through a carrots-and-sticks approach: positive outcomes are positively reinforced (rewarded) and negative outcomes are negatively reinforced (punished). Indeed, advancements in the quality of AI assistants, which can now perform complex "reasoning" (e.g., ChatGPT models ending in "o") have been achieved through reinforcement learning. In order to perform reinforcement learning, user feedback is required on what constitutes "good" outcomes that shall be reinforced. In order to scale this approach, typically one AI model is trained (based on user feedback) to be the judge of the outcomes provided by another AI model, which is to be improved. Evidently, the more user feedback is available, the better the judging model is aligned with the users' underlying preferences, and hence the better is the quality of the outcomes provided by the AI assistant.

Further, in order to turn an LLM into an assistant that responds to user requests, large conversational datasets – which serve as a model for conversations between humans and chatbots or agents – are required. While there are open-source conversational datasets available (e.g., derived from movies), creating and curating new/proprietary conversational datasets is an important source of competitive advantage for GenAI. For example, available conversational datasets may not be representative of the specific domain or use case that the model is being trained for. Thus, being able to derive conversational datasets from the interactions between users and the agent can prove a valuable source for improving the agent's quality.

Taken together, it is therefore reasonable to think that data-driven network effects are very relevant in the context of agentic AI. This is even more so true for agentic AI systems that serve as personal assistants, fulfilling tasks that a human personal assistant would, such as organizing travel itineraries. In this context, the AI system, like a human assistant, needs to know the idiosyncratic tastes and preferences of its user. User feedback data is thus a vital input to agentic AI, more so than for non-agentic AI.

Next to learning, another dimension along which agentic AI systems differ from AI assistants is its ability to perceive and interact with its environment. AI assistants, as we know them, typically run in a browser or app window, and interact with the user predominantly through text, voice or image input and output. Agentic AI systems are characterised by their ability to sense and interpret environmental data (e.g., through sensors), and to affect and act in the environment through so-called "tools". Tools are typically software components that enable the AI to perform certain actions. This is not limited to the virtual environment (e.g., moving the mouse, taking a

screenshot); it can also result in actions in the real word (e.g., turning the lights on, commanding a robot, booking a flight). In order to perceive and interact with its environment, agentic AI needs access to a larger array of inputs and to tools. In practice, this requires deep integration of the agentic AI system with the OS of the device on which the AI agent runs. For example, AI agents may require access to the device's GPS location, calendar, maps and browser to perform the tasks of an automated travel booking. Further, AI agents require access to tools provided by external content and service providers, such as mobility apps, search engines or payment providers in order to fulfil their tasks, such as searching, booking and paying a travel itinerary. These tools need not necessarily be downloaded and reside as software applications on the phone. Instead, the information exchange between the AI agent and the application could be handled purely through APIs.

For now, the prototypes of agentic AI systems released by OpenAI, Google, Anthropic and Perplexity lack deep integration with the OS and therefore rely on mimicking user input in the browser. Browsers in itself are already benefitting from deeper OS integration than other apps, and by their nature, can access all web content. Browsers thus open up a window for interacting with the environment that does not require further integration or permissions by the OS. In the future, however, integration must and will go deeper. As we highlight in more detail below (Section IV.A.5), agentic AI and the OS may even integrate to the extent that they cannot be separated anymore.

Another key dimension of the "agenticness" of AI systems is the ability to plan and reason (based on their perceived inputs). Planning and reasoning capabilities are built on the foundation of AI models. These come in two types. The first type is the general-purpose foundation model, such as those pre-trained transformer models invented at Google and popularised by OpenAI. Foundation models (which include LLMs) are large-scale, general-purpose AI models trained on broad and diverse datasets – practically all information accessible on the web. They are designed to perform well across many different tasks without needing major changes to their architecture. Training such foundation models is very cost-intensive and the market for the provision of such foundation models may end up concentrated. The second type is the specialised model, which is finetuned or built specifically for a narrow task or domain. Specialised models can be built from scratch but are often based on foundation models optimised for performance in a specific area. An AI agent will typically only make use of one foundation model but resort to several specialised models for specific tasks, such as shopping online, or scheduling events. In fact, some note that in the future an ecosystem of AI agents may emerge, where more general-purpose agents break up certain tasks and complex workflows into smaller subtasks, and then cooperate with or rely on more specialised agents to fulfil some of the (sub-)tasks.

While both foundation models and specialised models can be provided for the AI agent, fully or in part by third parties, the core of an AI agent is the agentic framework, which puts the various building blocks together. This is what makes AI "agentic": going beyond passive chat to take actions and plan sequences toward goals autonomously. This includes task planning, memory and

state management, goal prioritisation and execution, and deciding which models to use for which task. It also includes the integration code for tool access, i.e., a program that translates the AI agent's abstract tool requests (like "move mouse" or "take screenshot") into actual operations in the virtual environment, and the "agent loop", i.e., a program that handles communication between the agent and the environment, sending actions to the environment and returning the results (screenshots, command outputs) back to the agent.

Given the nascent state of agentic AI, there are few industry standards on what comprises an agentic framework and how agents can interact with the environment, but this is likely to undergo major developments in the future. Ultimately, the framework is the part that defines the agent, and the basis on which agents compete as it cannot be sourced from a third-party.

AI agents require access to computing resources, both for training and running (inference). For AI agents running on mobile devices, mobile AI chips capable of running AI models at the edge (i.e., on-device) are important to increase user experience, as they reduce latency and dependency on connectivity compared to cloud-based solutions. Furthermore, AI computing "on chip" can improve privacy as sensitive data does not have to leave the device. Currently not all advanced AI models can be run on-device, although the capabilities of on-device or edge AI are rapidly improving. Thus, AI agents are likely to employ hybrid solutions, where some tasks are performed by using AI on-device and others require access to AI models deployed in the cloud.

Developers of AI agents also require computing resources (in the cloud) to build agentic frameworks and/or train AI models. However, training AI models, especially foundation models, involves a complex value chain and resources of its own – one that relies heavily on existing gatekeepers in the digital economy. The CMA highlights that current gatekeepers under the DMA are active (sometimes only through partnerships) in all parts of the value chain, ranging from access to computing resources and data in order to train models, over development of foundation models, to model deployment in specific application areas such as search, mobile ecosystems or productivity software. This deep involvement of current gatekeepers has led commentators to conclude that there is "no AI without Big Tech" and that especially new startups have a high dependency on existing digital gatekeepers.

In our view, AI agents can in principle be built from modules, i.e., individual components that can be developed and supplied independently, and which work together through well-defined interfaces. Such modules can be (i) foundation models, (ii) specialized models, and (iii) the agentic framework, but also the (iv) hardware (e.g., device, sensors, chips), (v) operating system on which the AI agent runs, and (vi) cloud computing resources. Each of these modules could be supplied by a different provider, allowing a user to mix-and-match these modules in order to build a tailor-made agentic AI system. In this vision, competition for agentic AI would revolve around the agentic frameworks, rather than, e.g., foundation models. However, it seems that it is especially the providers of foundation models that are also pursuing the development and advancement of AI agents. Thus, while we have described the key components of an agentic AI system in a modular fashion, vertical integration between several of these key components is likely in practice. Vertical

integration is likely to occur between the agentic framework and AI (foundation) models, between agentic framework and OS and/or hardware providers, or both. As the market is still evolving, the extent of vertical integration in the long-run is yet to be determined. The experience of mobile ecosystems has shown that full-stack vertical integration, ranging from chips/hardware over OS and software – as in the case of Apple – is feasible, despite complex value chains.

3. POTENTIAL COMPETITION CONCERNS

The way in which AI agents are developed and deployed can raise competition concerns.

Competition authorities studying (Gen)AI have mostly focused on the development of foundation models, not on the specific issues raised by AI agents. This is not a problem, as most potential issues are known concerns manifesting themselves in a new setting. One can divide these concerns in two broad categories. First, a firm may promote its own AI agent and/or exclude that of competitors. Doing so requires market power either upstream (over one of AI agents' key inputs) or downstream (over important distribution channels). Second, a firm can use the AI agent itself to engage in anticompetitive conduct, e.g., by using it to self-preference products further downstream (e.g., specific apps). In order to do so, the firm operating the AI agent must have substantial power in the market in question – a market that may be broad (e.g., AI agents) but could also be narrower when consumers are locked in. The latter situation may arise when AI agents are integrated in another product (e.g., mobile devices) or if the costs of switching between agents is high (e.g., because their years of built-up memory are highly relevant).

In this section, we take a closer look at the two categories of potential issues: foreclosure of an AI agent (Section III.A) and foreclosure through an AI agent (Section III.B). Note that these two concerns roughly map onto the respective goals of the DMA. Preventing the foreclosure of AI agents is a matter of contestability, as this new technology has the potential to disrupt existing gatekeepers in a variety of core platform services, from search to OSs. But the providers of AI agents can also become gatekeepers in their own right. In that scenario, combatting foreclosure through AI agents is a way to ensure fairness in markets dependent on that core platform service (insofar as AI agents qualify as such – see below, Section IV).

A. Foreclosure of an AI Agent

Foreclosure of an AI agent can proceed via control over key inputs or important distribution channels.

1. Input foreclosure

An AI agent has a number of essential building blocks (see above, Section II). It is built on a foundation model that powers it. That foundation model is, in turn, trained by an expert workforce on a large dataset and with significant computing power (provided by chips directly or indirectly via a cloud service). Apart from the foundation model, the effective operation of an AI agent

requires retrieving data from the environment (e.g., by using the camera or viewing your screen) and from other services (e.g., your calendar, e-mails, etc.) as well as computing power for inference, which can be provided on-device or through the cloud. The essential nature of each of these inputs depends on the evolution of the market but when an input does become essential, its owner can use it to foreclose other AI agent developers.

AI agent developers currently have a choice of leading foundation models. While developing one's own foundation model is an expensive undertaking, developers can access those of others either by relying on an open-source model (e.g., Meta's Llama) or by relying on a more closed model that is made available through API calls (e.g., OpenAI's GPT).48 There is a risk that model providers cut off or degrade such access, e.g., as part of an 'open early, closed late' strategy.49

The alternative to relying on a third-party foundation model, i.e., for agent providers to develop their own, may be difficult if – even abstracting from the high cost – the key inputs are not

available.50 Compute, both through chips51 or cloud services52, is concentrated but remains accessible in Europe. The availability of training data may be a larger stumbling block. While models are primarily trained on data crawled from the open web, websites can – and increasingly do – block crawlers. Moreover, as the web has been crawled, training a competitive foundation model may increasingly require access to proprietary datasets. Some firms developing AI agents already own such datasets (e.g., Google's YouTube) and are likely unwilling to share them. An alternative is to license data from willing providers (e.g., social networks such as Reddit). However, if licensing agreements become widespread and are conditioned on exclusivity with a single developer, other developers may not be able to get the qualitative data required to train their foundation model. In addition to compute and training data, specialized talent is also limited and hence expensive.

Once powered by a foundation model, an AI agent needs a carrier. Unless agent-specific devices (e.g., the Rabbit's R1 or an in-the-works OpenAI device) break through, that carrier is likely to be the user's desktop or mobile device. Access to a carrier straddles the line between input and distribution: a carrier is necessary to be put in front of users (i.e., to be distributed, see further under "distribution foreclosure"), but an agent developer also needs access to its specific hardware and software features (i.e., input, see the next paragraph). Arrangements between agent developers and carriers may be complicated if the device manufacturer or OS developer either has its own agent or has an (exclusive) partnership with another agent developer. In those cases, it may not want to give (the required degree of) access to the independent agent developer.

Even if the independent AI agent can be installed on the device/ OS in question, it still needs access to hardware and software features controlled by the device/OS provider and by other parties (e.g., developers of specific apps). Say a user asks their agent to buy the sweater that is being advertised on a billboard and to have it delivered when they are home. The agent can only seamlessly do so when it gets access to (i) the mobile device's camera to snap the sweater; (ii) the user's shopping app or browser to order the sweater; (iii) the user's mobile payment mechanism to pay for the sweater; and (iv) the calendar app to check when the user is available for delivery. The incentives

of each third party to grant interoperability are complex, even when they do not compete against the AI agent provider in question. But wide access to hardware, software and a user's private information is crucial for an agent to effectively carry out tasks. Competition law sets a high bar for access to hardware and software (i.e., interoperability), though it has shifted somewhat in the Android Auto judgment. If the owner of the hardware/software has reserved it

On the (continued) availability of key inputs, see also Zach Meyers and Marc Bourreau, 'What policy interventions for a competitive AI sector?' (CERRE Report, July 2025).

The most high-performance chips are provided by a single firm, Nvidia, which is being scrutinized by competition authorities, see, e.g., 'French competition authority confirms investigation into Nvidia' (Reuters, 15 July 2024) https://www.reuters.com/technology/french-competition-authority-confirms-investigation-into-nvidia-2024-07-15/. 52 The cloud market oligopolistic, with the majority in the hands of Amazon, Microsoft and Google, see Ofcom, 'Cloud services market study' (Final Report, October 2023).

for its own use, the stringent refusal to supply test of indispensability (see below, Section III.B) applies; if that hardware/software is developed with a view to enabling third parties (e.g., is an open platform), a lower standard requiring simple anticompetitive effects applies.

2. Distribution foreclosure

As noted above, an AI agent needs a carrier. Complications can arise if the provider of the carrier provides its own agent or has a preferred agent (e.g., due to an exclusive/revenue-sharing arrangement with another agent developer). In that case, the carrier may want to give prominent placement, via pre-installation or increased visibility, to its (preferred) agent. By the same token, the carrier may want to make it more difficult for independent agent providers to reach their users. This can result in scenarios of distribution foreclosure, which have given rise to several antitrust infringement decisions in the digital economy over the past years. Let us take inspiration from them to get more specific on potential competition concerns.

First, an AI agent could be pre-installed (and made the default) on a popular device/OS, providing it with wide distribution. In previous cases, the pre-installation by an OS provider of their own app has been found abusive. Think, for example, of the pre-installation of Windows Media Player and Internet Explorer on Windows OS in, respectively, Microsoft I and Microsoft II.62 In Google Android, the issue was also one of pre-installation, though the Commission saw the tie in question differently: not between the OS (Android) and a specific app (e.g., Google Search) but rather between different apps (including Google Play, Chrome and Search). In that same case, another abuse was found in the agreements between Google and device manufacturers that provided the latter with a share of the revenue if they (exclusively) pre-installed Google's services, in particular its search engine.

Second, an AI agent could be integrated into a popular service, which may also guarantee its broad availability. The key case on integrating one service into another is Google Shopping, which turned on the integration of Google's comparison shopping service in its search engine. In particular,

Google Shopping was given prominence by being shown on top of the search results with enhanced graphical features. Facebook Marketplace concerned an even more relevant scenario, in which Meta integrated Facebook Marketplace into its social network. Legally, Google's conduct was qualified as "favouring" (or "self-preferencing") while Meta's conduct was conceived as tying.

Whatever the exact legal qualification, the pre-installation or integration of a specific app/service can become anticompetitive when it provides said app/service with a significant competitive advantage that cannot be offset by relying on other distribution channels.67 There is thus always a question of alternatives (other devices/OSs/services) through which the independent agent developer can reach users. If the dominant firm's distribution method makes its (preferred) agent ubiquitous, the threshold can be met. In addition, if the dominant provider of the

device/OS/service in question not only gives prominence to its (preferred) agent but also actively excludes that of competitors (e.g., makes the competing agent unfindable), anticompetitive conduct becomes easier to recognise.

B. Foreclosure through an AI Agent

AI agents could foreclose because they make recommendations to the user and even take decisions on users' behalf. In carrying out tasks for their users, they constantly make recommendations ("you're asking me to buy you a new Aran jumper; how do you like this option from Aran Sweater Market?"). As they become more autonomous and learn their users' preferences, they make more of the decisions independently, which cuts out the recommendation stage altogether. This is likely to happen first for low-cost decisions, such as where to order the margherita pizza the user asked for (via DoorDash or the Domino's website?), but may gradually expand to more significant decisions.

AI agents' recommendation and decision power can be harnessed to prop up its developer's other services through leveraging and tying. If the agent developer also happens to run a marketplace or food-delivery service, it may send requests for sweater and pizza purchases in that direction. This was a concern that scholars already voiced with regard to voice assistants, which were to become "gatekeepers for consumption" that could "hyper-nudge" users towards their preferred (vertically integrated) services.69 While such anticompetitive concerns were not borne out for voice assistants, they seem more realistic when it comes to AI agents.

Precedent on such anticompetitive conduct is found in the aforementioned Google Search decision, in which the Commission found Google's channelling of demand towards its own specialised shopping service abusive. The determinative question, in such self-preferencing cases, is whether the search engine/AI agent in question is an important source of traffic for the downstream service and whether the diverted traffic cannot effectively be replaced.

There is, however, an important difference between AI agents and search engines or other platforms that currently channel demand such as online marketplaces. When a search engine or marketplace puts its own product on top of the results and competing ones further down, the selfpreferencing is somewhat transparent. Users may still be under the (false) impression that the

platform's own product is in first place because it is the best product, but at least they can see that the platform is pushing that product. When an AI agent makes decision for its user, the selection process may happen entirely "behind the scenes" (unless the user prompts the agent to explain its actions). The self-preferencing can thus be more hidden and therefore more effective as it cannot be countered by users not clicking the first result; they would have to give specific instructions to use this or that service, but many users may leave that selection work to the agent. The risk of autonomous decisions that benefit the agent's provider but not the user is especially large when users do not pay for the agent, whose provider thus has to be remunerated by other parties (e.g., those to which it channels demand). The effects of self-preferencing may be greater when an AI agent is pre-installed or integrated in a device/service, in which case users – even when aware of bias – have a harder time "escaping" it.

Additional competitive concerns may arise when a "winner takes all" dynamic asserts itself and a single AI agent becomes inescapable. Such a dynamic may be spurred by indirect network effects: if third-party apps/services need to be reconfigured to be compatible with an AI agent, developers may only want to do so for a few agents, which can have a concentrating effect. If the market tips towards a certain AI agent and its developer then refuses to deal with, for example,

an OS provider that would like to integrate the agent, this can harm competition. An incentive to refuse in such scenarios only exists if the agent developer stands to gain, for example because its own OS is set to win market share as a consequence of the refusal. In such situations, the refusal to supply doctrine may apply. It sets a high bar: the AI agent must be indispensable (meaning there are no actual or potential substitutes) and the refusal must eliminate all competition (on the part of the firm requesting access).

Finally, as users rely more on AI agents, and the agents learns its users' preferences more and more over time, this bears the risk of data-induced lock-in effects. Users may be reluctant to switch to a provider of a rival agent if they cannot take their data and interaction history with them. Data-enabled learning can occur both within a user, but also across users, and can constitute a significant source of competitive advantage over other providers that do not have access to similar user data at scale. Thus, data-driven network effects may exacerbate the "winner takes all" dynamic and, especially if coupled with foreclosure of data access by the dominant AI agent provider, further hamper rival agent providers' ability to compete.

4. IS AGENTIC AI WITHIN THE DMA'S SCOPE?

As the AI agent market is only nascent, the exact risk and magnitude of any competition concerns remain uncertain. In the previous section, we have sketched potential concerns based on the experience in other digital markets, accounting for the specific features of AI agents. In case anticompetitive effects do manifest themselves, we gave a short indication of the extent to which competition law may apply. However, in view of competition law's limitations in digital markets, EU lawmakers adopted the DMA. Insofar as the DMA applies, it becomes – due to its procedural

efficiency – the Commission's first choice to tackle any competitive issues that may arise. Moreover, the DMA is meant to provide a "holistic regulatory response" to problems posed by gatekeepers (see above, Section I), which are likely to affect the AI agent market. In this section, we therefore ask whether agentic AI falls within the DMA's scope.

The DMA applies to gatekeepers, i.e., providers of core platform services (CPSs) that have been designated with gatekeeper status.75 Thus, in order to be regulated under the DMA, a service must first meet the definition of at least one of the ten CPSs. With our preceding description of AI agents in mind (Section II), we can immediately exclude "online social networking services", "video-sharing platform services", "number-independent interpersonal communications services", "online advertising services" and "cloud computing services" as conceivably fitting categories for the designation of AI agents. This leaves us with the following options for CPS categories that may lend themselves for interpretation to the context of agentic AI:

online intermediation services;

online search engines;

web browsers; • virtual assistants;

OSs.

We start by discussing whether these existing CPS categories accommodate AI agents (Section IV.A). Then, we consider the potential gatekeeper status of AI agents qualified as a CPS (Section IV.B), before providing a final analysis of the DMA's scope with regard to AI agents (Section IV.C).

It is worth nothing, before we start, that the DMA already has grasp on AI services even without their designation. In particular, the source of foreclosure of AI agents (Section III.A) is likely to be found in CPSs with existing gatekeepers, first and foremost OSs. Many of the

obligations – e.g., relating to interoperability – apply irrespective of any AI agent designation, though there are exceptions – e.g., the choice screen obligation (see below, Section V). The source of foreclosure may also be found in other CPSs. An example is provided by Article 5(2) DMA, which prohibits gatekeepers from cross-using personal data collected from CPSs in other services (including non-gatekeeper ones), unless the user consents. This would appear to bar a gatekeeper like Meta from using data collected on its designated social networks to train its (non-designated) AI services.

A. Designation as Core Platform Service

As there are five CPSs which AI agents may qualify as, a preliminary question is whether a service needs to be designated as a single CPS or whether a combinatorial designation is possible. When setting thresholds for designation, the DMA does refer to CPSs in the singular. The issue came up in the designation of TikTok. ByteDance, TikTok's owner, argued that it should be designated as a video-sharing service; instead, the Commission designated it as a social network. There were

valid arguments for either position: video-sharing is TikTok's primary purpose and core function, while its breadth of other features brings it closer to a social network. Neither ByteDance nor the Commission explored a dual designation and, as ByteDance did not dispute its qualification on appeal, we did not hear the General Court's opinion. Hence, also in view of the practical difficulty combinatorial designation would cause (e.g., when counting users), it seems safe to assume a service can currently only be designated in one CPS category. This makes the finding the right category all the more important, especially as the obligations also vary per category (see below, Section V).

1. Online Intermediation Service

Online intermediation services are defined with reference to the P2B Regulation as (i) information society services (ii) that allow business users to offer goods or services to consumers with a view to facilitating direct transactions between those business users and consumers, and (iii) that are provided to business users on the basis of contractual relationships between the provider of those services and business users which offer goods or services to consumers.

The definition emphasises intermediation between consumers, on one side, and business users (offering products or services to consumers via the online intermediation service), on the other.

While an AI agent may be employed as a tool to discover and transact with businesses (e.g., an AI shopping agent), typically there is no contractual relationship between the business user (e.g., third-party shops) and the AI agent. Rather, the agent connects to applications or, indeed, online intermediation services (e.g., online marketplaces); alternatively, it surfs the web to complete the purchase, as in the case of Amazon's agentic "Buy for Me" feature. The applications or websites visited by the agent may have business users, but the agent itself does not (directly).

The provider of the AI agent could have contractual relations with some of the business users that it may interact with, but the nature of the AI agent is such that – following a user's request – it would autonomously search for relevant matches and perform actions independent of preexisting contractual relations with the parties that it interacts with. The presence of contractual relations with business users is not a defining feature of AI agents.

In conclusion, as AI agents do not directly host business users, nor necessarily have contractual relations with them, they generally do not qualify as an online intermediation service in the sense of the DMA.

2. Online Search Engine

The DMA defines online search engines also with reference to the P2B regulation as "digital service[s] that allow … users to input queries in order to perform searches of, in principle, all websites … on the basis of a query on any subject in the form of a keyword, voice request, phrase or other input, and returns results in any format in which information related to the requested content can be found".

This definition contains some of the core features of an AI agent. Indeed, an integral part of AI agents is that they are able to autonomously retrieve external information, such as those provided by, in principle, "all websites". This is already a prevalent feature of AI assistants, which resort to new information collected from the web in order to combat hallucination and to overcome the knowledge cut-off inherent to the training data. Web search has therefore become an integral part of modern LLM-based systems. However, the functionality of AI agents is of course not limited to searching the web and goes beyond information retrieval – a point to which we return below.

One could also argue that AI agents, particularly those embedded in personal devices, perform their actions in response to a "query on any subject", i.e., a user request. The definition of search engines allows for a wide range of possibilities in which this request can be formulated, including text, voice or "other input". Thereby, "other input" can be interpreted as a catch-all term that possibly opens the door for all additional conceivable options in which a human may interact with technology, such as AI agents, including, but not limited to gestures, facial expressions or even through neural interfaces.

Finally, the format in which the search engine must provide the results in response to a user request is also defined in an open-ended fashion: it can be in "any format". Therefore, the form in which the AI agent responds and delivers "results" to the users' request does not disqualify it from being a search engine.

For example, a user that formulates a request to the AI agent to "buy tickets the cheapest tickets for the next concert of Shania Twain", which requires the AI agent to "search" for cheap tickets on, in principle, all websites (ticket sellers, private listings, online marketplaces, online communities, etc.) and returns results in form of a valid (electronic) ticket, may fit the definition of "search engine" under the DMA. The AI Act also recognizes how "AI systems may be used to provide online search engines".

However, as noted, an AI agent is not limited to search functionality. It also caters to requests that do not necessarily involve the search of all websites. Further, even though AI agents are likely

to encompass a search function, and in turn all major search engines tightly integrate with AI, the integration between a search engine and the AI agent can vary. The few providers of AI agents that have their own search engine can rely on it, providing for a tight integration (e.g., Microsoft's Bing AI, Google's Gemini). Other providers have to rely on one of the main search engines (e.g., OpenAI's ChatGPT taps into Bing). In either case, the question is whether the AI agent is a search engine or simply makes use of a search engine.

There is a fine line between the two, which is getting even more blurred as AI agent providers launch ever more search-like products (e.g., OpenAI's SearchGPT). There are therefore valid arguments to qualify and disqualify AI agents as search engines. The DMA's text provides an argument in favour of doing so. But as long as standalone search engines exist and are distinguishable from agentic AI systems, one may argue that AI agents are by nature different and simply make use of a search engine. Then again, given the rapid evolution of search engines in

response to AI, in a not-so-distant future, standalone search engines as we know them may no longer exist or at least no longer play a significant role. Instead, search is likely to become a commodity fully embedded in AI-based consumer-facing services, such as AI assistants or indeed agents. In this case, search – as envisioned by the DMA – could only be performed through these AI-services, and thus AI agents may well be considered (at least in part) a search engine.

3. Web Browser

The DMA defines a web browser as "a software application that enables end users to access and interact with web content hosted on servers that are connected to networks such as the Internet, including standalone web browsers as well as web browsers integrated or embedded in software or similar".

Reminiscent of our discussion of AI agents as search engines, AI agents may well embed a web browser as an integral component. The functionality of the current prototypes of AI agents all heavily rely on a web browser to perform their tasks. Moreover, the definition of a web browser in the DMA is very broad, similar to that of search engines, encompassing virtually any software that allows to connect to "web content", whether it be on the (public) Internet or any other servers connected to (private) networks. Similar as in the case of search engines, however, the functionality of AI agents clearly goes beyond what is commonly considered a web browser today. Again, the central question is therefore whether an AI agent is a web browser or just makes use of a web browser.

Looking ahead, it is likely that AI agents will fundamentally challenge the notion of standalone browsers as we know them. Searching and browsing the web will be done for us, rather than by us. To facilitate this task, firms like Perplexity and OpenAI have already developed their own, agent-ready version of a browser. We expect the distinction between AI agents and browsers to become increasingly blurred, just like in the case of search engines.

4. Virtual Assistant

A CPS category particularly reminiscent of AI agents is that of "virtual assistant".90 The DMA defines a virtual assistant as "a software that can process demands, tasks or questions, including those based on audio, visual, written input, gestures or motions, and that, based on those demands, tasks or questions, provides access to other services or controls connected physical devices".91 Whether AI agents are captured by the DMA as virtual assistants depends in part on the relative weight one gives to the legislative text versus the underlying intent at the time of adoption.

Starting with the text, the definition of virtual assistant is – as that of search engines and browsers – very broad. In its plain meaning, the definition would capture AI agents, whose essence is to process tasks or questions based on a variety of cues (voice, writing, gestures, etc.). Semantically, some make a distinction between assistants and agents. IBM, for example, puts it as follows: "AI assistants are reactive, performing tasks at your request. AI agents are proactive, working autonomously to achieve a specific goal by any means at their disposal." This semantic distinction

is unlikely to matter, however, as qualification is based on substance – not labels. Moreover, the distinction obscures that assistants and agents are situated on a spectrum, with current LLM-based services moving from the assistant to the agent side. Based on the substance of the definition, AI assistants and the agents they are evolving into may thus qualify as virtual assistants. In the more distant future, as agents reach a high degree of autonomy, the definition may become less fitting as there may no longer be (explicit) "demands, tasks or questions" (prompts) given by the user; the AI agent will proactively take action based on goals previously set by the user or even just inferred by the agent.

The legislative history somewhat complicates this picture. Virtual assistants were not included in the original DMA proposal but were later added by lawmakers. At the time, the Commission was carrying out a sector inquiry into the consumer Internet of Things (IoT), with a focus on voice assistants such as Amazon's Alexa. While there were competitive concerns, e.g., around interoperability, the Commission did not launch a subsequent antitrust investigation. Concerns did find their way into the DMA's legislative process, where the European Parliament picked them up. A dive into the DMA negotiating process shows how the Parliament at one point considered including "voice assistants" as a CPS but later broadened the CPS to "virtual assistants". In line with this amendment, lawmakers included a definition of virtual assistants that is much broader than the voice assistants scrutinised at the time of the DMA's adoption. The reasons for this change are, however, not part of the public record.

Nevertheless, there remain some signs in the DMA itself that lawmakers (initially) had traditional voice assistants – not AI agents – in mind during the drafting process. First, the "virtual assistant" definition refers to the control of connected physical devices, which is a typical feature of Alexa-type voice assistants. Second, the DMA defines business users of a virtual assistant as "developers who offered at least one virtual assistant software application or a functionality to

91 DMA, art 2(12).

make an existing software application accessible through the virtual assistant". This seems to hint at the "skills" developed for specific voice assistants like Alexa. The idea of AI agents, by contrast, is that they can autonomously access various apps, which does not necessarily require apps to be developed specifically for the AI agent. We are, however, seeing the first signs of agentspecific API frameworks, and of partnerships between agent and app developers.

In conclusion, the DMA's "virtual assistant" category provides in our opinion a rather close fit for AI agents, at least in the near to mid-term future. The definition certainly accommodates them, even if it is not fully clear what the legislative intent was when adding this specific CPS to the DMA. In the longer-term, and in view of policy considerations such as legal and investment certainty, a legislative rather than interpretative route may be preferable. We return to this in our concluding analysis (Section IV.C).

5. Operating System

In the DMA, an OS is defined as "a system software that controls the basic functions of the hardware or software and enables software applications to run on it". At present, such a qualification seems far-fetched: AI agents are, at most, an app or website accessed via an OS; they are not themselves OSs. Looking to the future, however, that might change.

For an AI agent to operate effectively, it needs access to a device's hardware and software (see above, Section II). Over time, it may not just get access but actual control. Already, firms are deeply embedding these agents in their OSs. Apple Intelligence, for example, is said to be "integrated into the core of your iPhone, iPad, and Mac". On the other side, developers can use Apple's App Intents framework to "deeply integrate" their app's functions and content with Apple Intelligence. As a result, Apple Intelligence can call upon the app for specific tasks without the user having to navigate there. However, while the AI is invoking the apps, the apps are not yet "running on" the AI agent.

In the medium to long term, AI agents may not just run apps but also take over control over the device's hardware and software from the OS – in effect, become the OS. The first signs of this shift are discernible in OpenAI's introduction of apps within ChatGPT, which hints at an ambition to become an OS in its own right. While socio-technological developments are likely to bring AI agents closer to OSs over time, however, they are unlikely to qualify as OSs in the sense of the DMA in the near term.

B. Determination of Gatekeeper Status

After settling on a fitting CPS category, the AI agent provider in question would need to achieve gatekeeper status before being subject to the DMA's obligations. Gatekeeper status under the DMA is conditioned on three qualitative criteria: (a) having a significant impact on the internal market; (b) serving as important gateway for business users to reach end-users; and (c) enjoying

an entrenched and durable position. These criteria are presumed to be fulfilled when a firm meets certain quantitative criteria, respectively:

achieving EU revenue of at least €7,5 billion in each of the last three financial years or having an average market capitalisation/fair market value of at least €75 billion in the last financial year, and providing the same CPS in at least three Member States;

providing a CPS that has at least 45 million monthly active end-users and 10.000 yearly active business users in the EU in the last financial year; and

meeting the threshold of (b) in each of the last three financial years.

These thresholds raise some questions when it comes to AI agents. Criterion (a), on revenue/market cap (or fair market value) remains straightforward. The interpretation difficulty is centred on criterion (b), regarding the number of end- and business users. Identifying users that have engaged with the AI agent in the past month should pose no difficulty, but identifying business users does.

According to the DMA, a business user is "any natural or legal person acting in a commercial or professional capacity using core platform services for the purpose of or in the course of providing goods or services to end users". This definition seems to imply that a business user must actively choose the CPS to conduct business. The specific definition of business users of virtual assistants, discussed already above (under "Virtual Assistant"), also appears to require active conduct, as it covers "unique developers who offered at least one virtual assistant software application or a functionality to make an existing software application accessible through the virtual assistant".

In the context of AI agents, however, it is questionable whether the business user has actively chosen the AI agent to do business, or the other way around (see already above under "Online Intermediation Service"). Certainly, some business users are actively relying on AI agents via partnerships and/or API integration. But "computer use" models seek out business users without their active cooperation.

Business user identification thus poses a potential problem when designating AI agents as virtual assistants. One could get around this problem with a wider interpretation of business user, which considers any business called on by an AI agent to be a business user, given that this business is using a CPS "in the course of providing goods or services to end users". This is the approach taken by the DMA for search engines, where a business user is anyone "with business websites indexed by the online search engine". Clearly, the search engine is the party doing the indexing, whereas the business websites are passively being indexed (though they can take active steps to resist such indexing). One may therefore want to analogously interpret business users of virtual assistants as encompassing also those that are chosen by the AI agent to provide goods or services to end-users. As this requires quite an interpretative leap, however, we recommend adapting the business user definition (see below, Sections IV.C and VI).

Criterion (c) sets an entrenchment threshold, i.e., boasting the above number of business and end-users for three years. There is thus a time-lag built into the DMA, based on the idea that – even when certain CPSs are seeing (momentarily) high user numbers – they need to show a degree of stability before being designated with gatekeeper status and subjected to regulatory obligations. Any new technology, including AI agents, will thus see the designation of gatekeepers at the earliest three years after its widespread adoption. There are, however, two exceptions to this rule.

First, the Commission could designate an AI agent based on a market investigation, in which it establishes that the qualitative but not the quantitative gatekeeper criteria are met. It would then still have to show that, based on the DMA's criteria and barriers to entry, the provider is entrenched with respect to its AI agent. Given that AI agents are nascent and that the market is, for now, still heavily contested, this seems like a tall order in the near-term.

Second, the Commission could designate an AI agent provider as "emerging gatekeeper". Current entrenchment is not required in this case; rather, the question is whether the AI agent provider "will foreseeably enjoy [an entrenched and durable] position in the near future". This also requires a market investigation, and the Commission can only declare specific obligations applicable to the

emerging gatekeeper. At the time of writing, no CPS provider has been designated as emerging gatekeeper, and no AI agent seems to fit that bill for now.

5. Analysis

AI agents challenge the DMA's CPS categorisation. Of the five CPSs that appeared to provide a plausible categorisation (i.e., could not immediately be excluded), there are – at present – three categories that more or less fit. Online search engines and web browsers have both been defined broadly and, as such, could be interpreted to cover certain AI agents. The connection lies in the fact that AI agents widely make use of the web for the retrieval of information and to carry out tasks. Theoretically, it appears more correct to say that AI agents rely on search engines/browsers but do not constitute them. EU regulation, however, does not draw such fine a line – and that line in any case becomes increasingly difficult to draw as AI agents integrate more and more search/browsing functionality. Indeed, OpenAI has been reporting user numbers as a search engine for the purpose of the Digital Services Act.

For AI agents, a third CPS category – of virtual assistants – provides a cleaner fit. There are some unknowns in the legislative origins of this category but its inclusion and subsequent broadening from "voice" to "virtual" assistants suggests an openness to future developments. The DMA's text is a greater source of authority, and its definition of virtual assistants is – like that of search engines and web browsers – broad. The DMA's wide definitions also suggest an attempt to keep the instrument reasonably open to the rapid evolution of digital markets, in other words, to keep it future-proof. In line with that aim, the broad virtual assistant definition – despite textual remnants of the voice assistants that some lawmakers had in mind during the DMA's adoption – relatively comfortably accommodates AI agents.

This finding must be nuanced in two ways. First, AI agents are rapidly evolving. While the virtual assistant definition may cover the first generations of AI agents, that coverage can become tenuous over time. As AI agents attain ever more capabilities, autonomous action is likely to become more prevalent: agents may shift from processing user input/prompts ("demands, tasks or questions") to anticipating user "needs", at which point the virtual assistant definition becomes less fitting. Second, for several CPS categories including that of virtual assistant, the Annex's business user definition leads to issues, as AI agents do not necessarily or actively have them, even though this may also change over time if partnerships and API integration become more established practice.

The alternative to dealing with these issues through interpretation is legislative change. To account for the dynamism of the digital economy, lawmakers have built revision mechanisms into the DMA. Explicitly adding AI agents to the list of CPS (by changing or replacing the virtual assistant definition) would require the Commission to carry out a market investigation and to submit it, accompanied by a legislative proposal, to the European Parliament and the Council. This can be done as part of the DMA's three-yearly evaluation, of which the first one must be carried

out by May 2026. Adapting the definition of business users, in particular those of virtual assistants, is possible via a delegated act and is therefore less legislatively demanding.

In the choice between the two routes to capture AI agents under the DMA – interpretation and legislative change – there are different policy considerations at play. The main ones are effectiveness and legal certainty – two values that are central to the DMA. Interpretation provides for (immediate) effectiveness; legislative change comes with greater legal certainty, which in turns provides confidence to investors. In addition, legislative change confers greater legitimacy on the inclusion of AI agents within the DMA's scope, which may be particularly important in view of (i) the scant attention given to AI in developing the DMA, and (ii) the (political) pressure for more caution in creating regulatory burdens in the EU (from within the EU and from across the Atlantic).

The two routes are not mutually exclusive. If AI assistants gradually evolve into agents, and become entrenched and important gateways for business users to reach end-users, the Commission could designate them under the virtual assistant CPS category in the near to midterm. This ensures effectiveness and, given that the first generations of AI agents are close (enough) to the virtual assistants described in the DMA, legal certainty is not significantly imperilled. Before designation, however, it would be useful to adapt the business user definition via delegated act. If agents become ever more autonomous, which may be expected in the mid to long-term, legislative change would bring greater legal certainty and might even be needed to effectively capture AI agents. At that point, such change can be based on a clearer view of the trajectory of agentic AI, e.g., towards more system-level software. This two-step approach rhymes effectiveness with legal certainty.

Moving from CPS qualification to gatekeeper status, there are fewer difficulties on that front (apart from defining and consequently "counting" business users, discussed above). Designation is – by design – a matter of time, as the DMA does not intervene before significant adoption. Market investigations and the status of "emerging gatekeeper" providing a possibility to modestly speed up that timeline, but we see no immediate need to rely on them.

6. ARE THE DMA'S OBLIGATIONS MEANINGFUL IN THE CONTEXT OF AGENTIC AI?

In the preceding section we have determined that AI agents may conceivably be designated – provided they meet the criteria for (emerging) gatekeeper status – as either search engine, web browser, or virtual assistant, although the virtual assistant category is the most fitting of these. Over time, AI agents may come to resemble OSs, and they any case need some level of access to OSs to function. However, as none of these CPS categories were designed with the specific case of AI agents in mind, it is necessary to check whether the specific obligations that are attached to

these categories, individually or collectively, can still meaningfully mitigate the competition concerns associated with agentic AI (Section III).

While the DMA's obligations are sometimes seen as generally applicable to gatekeepers, many of them apply only to a specific CPS (or set of CPSs). As pointed out above (Section IV), this means that CPS qualification must not only be formally correct but also functionally relevant (i.e., in a category where obligations are relevant to the service in question). To underline this importance, we discuss the relevant obligations – as they pertain to AI agents – based on the CPS to which they apply: OS (Section V.A); OS, virtual assistant, and web browser (Section V.B); search engine (Section V.C); all CPSs (Section V.D). We do not discuss obligations that apply only to CPS categories in which AI agents cannot conceivably be designated.

A. Obligations for OSs

The DMA contains specific obligations that only apply to a designated OS, and not to any other CPS category, in particular in Article 6(3), paragraph 1, which demands that "the gatekeeper shall allow and technically enable end users to easily un-install any software applications on the operating system of the gatekeeper" unless those software applications are "essential", meaning that they cannot reasonably offered by a third-party provider on a standalone basis.

Provided that AI agents are not designated under the DMA as an OS, but the OS provider has pre-installed a (non-essential) AI agent, then this obligation (in conjunction with the interoperability requirements of Article 6(7) DMA discussed below) ensures that an alternative third-party AI agent can be installed on the OS. The provision is therefore meaningful in the context of preventing the foreclosure by an OS provider of a competitively supplied AI agent (see above, Section III.A). Note that it is not material for this provision whether the pre-installed software is from the gatekeeper or not. Thus, the provision should also apply to cases where an OS provider merely cooperates with the provider of an AI agent, e.g., in case of Apple cooperating with OpenAI for Apple Intelligence.

However, as agentic AI and the OS are likely to become more intertwined (see above, Section IV.A.5), a deep integration of the AI agent and the OS is required, so there will be debate around how "essential" the pre-installed AI agent is, and whether it could indeed be offered on a standalone basis. The latest advancements in OS development, both on mobile and desktop devices, already integrate AI support to a degree that well exceeds the integration of other typical third-party applications. If the embedding of AI agents into the OS is indeed deemed "essential", then this opens up the possibility to designate AI agents as an OS. In this case, there would be no uninstallation requirement according to Article 6(3), paragraph 1, but the gatekeeper providing the integrated OS/AI agent environment would yet need to provide equitable interoperability to third-party AI agent providers according to Article 6(7) DMA, since the OS-integrated AI agent is still a "software feature" provided by the gatekeeper. However, consumers may be more reluctant to install a third-party AI agent in case the pre-installed AI agents cannot be un-installed, which was precisely the legislative intent behind the uninstallation requirement expressed in the first paragraph of Article 6(3).

B. Obligations for OSs, Virtual Assistants, Web Browsers

In the DMA, virtual assistants are often put in a category different from "simple" applications.

In particular, they are lumped together with OSs and browsers, as the type of software that other

applications can run on or in. For system software, such as OSs, that function is clear. The inclusion of browsers in this category is understandable: the (US) Microsoft case was about how Microsoft saw a "middleware" threat in Netscape's browser, which could become a platform for applications in its own right and thus make Windows obsolescent. The inclusion of virtual assistants in this category is – or was – not immediately obvious. With its rudimentary "skills", for example, Amazon's Alexa did not look likely to replace OSs any time soon. But AI agents are more credible as a complement that may one day become a substitute for OSs (see above, Section IV.A.5). In that sense, the categorisation is deserved. Two obligations – one relating to default setting and choice screens, the other relating to interoperability – should be viewed in this light.

1. Default Settings and Choice Screens

Article 6(3), paragraph 2, DMA holds that: "The gatekeeper shall allow and technically enable end users to easily change default settings on the operating system, virtual assistant and web browser of the gatekeeper that direct or steer end users to products or services provided by the gatekeeper." Given that this provision is sandwiched between one that concerns only OSs (Article

6(3), paragraph 1, see Section V.A) and another that applies only to firms that are gatekeepers in OSs and in a select set of downstream services (Article 6(3), paragraph 2, part 2, see below), one might wonder about the scope of this provision. It should be interpreted as applying irrespective of whether the firm is designated in any other CPS category. Hence, should an AI agent be designated as either OS, virtual assistant or web browser, the gatekeeper needs to ensure that – insofar as the service defaults to other services of the gatekeeper (irrespective of whether those services have itself gatekeeper status) – it makes that default setting changeable. For example, if the gatekeeper provider of an AI agent would fulfil shopping requests through its (nongatekeeper) marketplace by default, it would have to allow users to change that setting.

The second part of this provision goes a step further, mandating choice screens for a select set of CPSs. Upon close reading, the selection is rather narrow: it only covers the designated "online search engine, virtual assistant or web browser to which the operating system of the gatekeeper directs or steers users by default, and the online search engine to which the virtual assistant and the web browser of the gatekeeper directs or steers users by default". Hence, the obligation applies in two scenarios:

A firm that is designated as gatekeeper in both the OS CPS category and in the search engine, virtual assistant or web browser CPS category must show a choice screen for the designated app category at the level of the OS. In line with this obligation, Apple shows an iOS browser choice screen (it is designated in the OS and browser CPS categories) but not an iOS search engine choice screen (it is designated in the OS but not the search CPS category, as Google provides the default search engine on iOS). The Commission has investigated the design of Apple's choice screen, but not the scope of its implementation

(i.e., the fact it is only shown for browsers, not search engines), confirming the obligation's coverage as interpreted above.

A firm designated as gatekeeper (only) in the virtual assistant or web browser CPS category needs to show a search engine choice screen. Alphabet, for example, shows a search engine choice screen on its Chrome browser (whatever the OS it is installed on).135

While choice screens have not always been as effective, this obligation can be meaningful in bringing AI agent choice – and thus in counteracting the foreclosure of AI agents by vertically integrated players (see above, Section III.A). In particular, users may make more effective use of that choice when the market has not fully settled (i.e., there is less bias towards the familiar). Whether an AI agent choice screen is shown, however, depends entirely on its designation. If it is designated as search engine, virtual assistant or web browser, OS gatekeepers also designated in that category must show a choice screen. This means that Microsoft (Windows), Apple (iOS) and Google (Android) would have to provide search/assistant/browser choice screens on their OS if their AI agent is designated in one of those three categories. If a firm's AI agent is designated as virtual assistant or web browser, it needs to show a search engine choice screen (but not a choice screen for other virtual assistants or browsers).

There are three problems with this obligation in an AI agent context. First, a potential doubling up could nullify certain obligations. If a firm's AI agent is designated as OS – a distant but not unthinkable possibility (see above, Section IV.A.5) – it might need to show a search/assistant/browser choice screen. However, an assistant choice screen is obviously without purpose in this scenario. Further, if the firm's (traditional) OS is also designated – i.e., it has two OS designations – it would not have to show any choice screen, as it is not designated at both the OS and the app level. This makes technical sense: if the agent essentially overlaps with the OS, you cannot just swap it out. In a future scenario with such deep integration, contestability is limited, but this may be outside of the DMA's scope: while it tries to bring contestability within OSs (e.g., through interoperability and alternative app distribution channels), it does not directly seek to bring contestability in the OS market itself. By contrast, a search or browser choice screen may be interesting here, especially as agents make great use of these adjacent services. Again, however, it may be technically difficult to swap out the search index the agent relies on.

Second, the more meaningful obligation that allows for an AI agent (not just a search engine) choice screen, is limited to vertically integrated firms that are gatekeepers at two levels: upstream (OS) and downstream (AI agent, whether designated as search engine, virtual assistant or web browser). This means that it does not apply when an OS gatekeeper contracts out its provision of an AI agent. Apple, for example, has partnered with OpenAI to provide its agent, and may yet partner with Google – as it has done for search. Insofar as those agents are not considered its own, it would escape any choice screen obligation on iOS even if the agents rise to the level of gatekeeper status.

Third, AI agents may communicate with each other, with one agent invoking another, possibly more specialised, AI agent to perform a task (see above, Section II). However, if an AI agent is

designated as a virtual assistant (or web browser), it does itself not have any obligation to offer a choice screen for which other AI agents it may invoke per default. Generally, this issue seems of lesser concern, however, as it may become very complex quickly if a user were always required to mix-and-match the ecosystem of AI agents she wants to use.

2. Interoperability

Article 6(7) DMA centres on interoperability and has two parts to it.

The first part of the provision obliges firms with a designated OS or virtual assistant to allow other hardware and service providers free and effective interoperability with hardware and software features as are available to services or hardware provided by the gatekeeper.

The second part of the provision obligates firms with a designated OS (so not virtual assistant) to allow business users and other providers of services (not hardware) that integrate with CPSs free and effective interoperability with OS, hardware and software features as are available to the gatekeeper when providing such services.

The goal of interoperability is that "all elements of hardware or software work with other hardware and software and with users in all the ways in which they are intended to function". The degree of interoperability to be provided is determined by a non-discrimination standard: other providers must receive the same interoperability affordances as the gatekeeper's own hardware and services – hence, the obligation is one of equal vertical interoperability. The provision is set up to prevent situations of input foreclosure, where an OS/virtual assistant gatekeeper uses interoperability to discriminate in favour of (self-preference) its own downstream services to the detriment of competitors (see above, Section III.A). There thus needs to be a degree of vertical integration (OS and downstream service) for the obligation to apply.

In the context of AI agents, the obligation can play a dual role. First, it is crucial in ensuring that the providers of third-party AI agents can compete on a level playing field with the OS provider's first-party AI agent (part 2). On Android, for example, the AI agents of third parties should get the same access as Google's own agent. There is a potential problem with building this obligation around vertical integration: if the OS provider contracts out the provision of an AI agent, does it still qualify as a "service provided by the gatekeeper"? Insofar as the OS gatekeeper's AI agent is provided in partnership with a third party, such as in the case of Apple and OpenAI (for Apple Intelligence), the obligation would still seem to apply. But if the OS gatekeeper pre-installs an AI agent provided entirely by a third party (in exchange for some payment), the equal interoperability obligation would seem not to apply.

Second, AI agents themselves can also be covered by the obligation if they are designated as either OS or virtual assistant (part 1). In that case, their providers do not benefit from interoperability but have to grant it to others. If the AI agent provider manufacturers a connected device to run, support or integrate with the agent, for example, it would have to grant competing hardware providers equal interoperability. As seen in a specification decision, such interoperability could relate to the display of notifications, proximity-triggered pairing, file transfer, wi-fi connection, near-field

communication control, and so on. It would also require the gatekeeper to institute a process to handle interoperability requests.

C. Obligations for Search Engines

When it comes to search engines, the obligations relate to self-preferencing, data sharing, and fair, reasonable and non-discriminatory (FRAND) terms. As we explain below, however, the selfpreferencing obligations may extend beyond search engines.

1. Self-preferencing in Ranking

Article 6(5) DMA prohibits self-preferencing: "The gatekeeper shall not treat more favourably, in ranking and related indexing and crawling, services and products offered by the gatekeeper itself than similar services or products of a third party." The wording of the prohibition – in particular, its reference to ranking, indexing and crawling – would seem to restrict it to search engines, which is probably due to the prohibition's origins in the Google Shopping case (see above, Section III). If an AI agent is designated as a search engine, this is not a problem. If an AI agent is designated in another CPS category, however, the provision's terminology would need some stretching in order to apply, though this does need not make the prohibition inapplicable.

For instance, an AI agent does not necessarily "rank" services and products in the sense that it provides a ranking to the user. Even though an agent may rank options for a user, it could also directly buy, subscribe or more generally consume on the user's behalf. The DMA, however, subscribes to a wide definition of "ranking":

the relative prominence given to goods or services offered through online intermediation services … or virtual assistants, or the relevance given to search results by online search engines, as presented, organised or communicated by the undertakings providing online intermediation services … virtual assistants or online search engines, irrespective of the technological means used for such presentation, organisation or communication and irrespective of whether only one result is presented or communicated.

This wide definition may have been adopted in view of Google's gradual evolution from a search engine (with different results – the proverbial "ten blue links") to an answer engine (with just one [prominent] result), or perhaps to cover virtual assistants, which tend to skip the ranking of different options and are explicitly included here. Given this definitional exercise, Article 6(5)'s self-preferencing prohibition would therefore seem to cover AI agents when designated as search engines or, indeed, in other CPS categories such as virtual assistants.

The second part of the provision, which obliges the gatekeeper to "apply transparent, fair and non-discriminatory conditions to such ranking" is also difficult to apply in the context of AI agents. Even if there is ranking going on (behind the scenes), it is a very challenging task to provide "transparency" about how an AI agent (typically based on a deep neural network) derives decisions and suggestions from its training and the available data. While explainable AI (XAI) is an active area of research, explanations need to be on a high level of abstraction from the underlying neural

network, and it is therefore doubtful whether they are really suitable to determine with sufficient legal certainty whether "fair and non-discriminatory" conditions have been applied.

Moreover, the ranking could be derived from several different components of the agentic AI system (see above, Section II), such as the (i) foundation model, (ii) a specialised model, and the (iii) the agentic framework. In our view, the final ranking can happen in each one or a combination of these three components. For example, the agentic framework could just rely on the ranking done by the underlying foundation model, or it could refine a ranking input received from a specialised model (which again used a foundation model as input). These are just two of the possibilities, which reinforces our point that it will be challenging to provide "transparency" and how exactly the ranking was derived, especially if this includes several components of the AI

system, which may even be supplied by different entities. The problem of this provision of the DMA is thus one of implementation and enforceability.

2. Data Sharing

A second obligation applicable to search engines is found in Article 6(11) DMA: "The gatekeeper shall provide to any third-party undertaking providing online search engines, at its request, with access on fair, reasonable and non-discriminatory terms to ranking, query, click and view data in relation to free and paid search generated by end users on its online search engines." More clearly than for Article 6(5) DMA, the addressees of this search data sharing obligation are limited to gatekeeper search engine providers. In moving from traditional search engines to AI agents, however, the scope of the covered data loses some clarity. "Query" data is still straightforward enough, as AI agents require input, even if it may come in a wider variety of forms (text, voice, visual…). With regard to "ranking", we pointed out the issue above. "Click and view" data, finally, refers to the typical follow-up actions on the search results page. Interactions with an AI agent may but certainly need not result in a click – the whole point of moving from search to answer engines is to obviate the need for such a click. Similarly, the user of an AI agent may not "view" any results at all; the action may simply be carried out (or the answer might be "heard" rather than viewed). This interpretation difficulty can be resolved by focussing on click and view data at one level of abstraction – so on follow-up actions generally rather than on specific methods of navigating the CPS – but even then, there is still the practical difficult of sharing certain data (e.g., ranking data).

Finally, the question is how useful the shared data would be for contestability in the market for AI agents. Like in the case of search engines, the answer to this question for AI agents likely also depends to a large extent on the details, such as the granularity and extent of the data provided, and the degree of anonymisation that is done. Generally, as we have discussed in Section II, AI agents can improve and learn from user feedback, and thus, access to query, click and view data seems to also be, in principle, meaningful in the context of agentic AI. However, given the broader scope and more complex operations that agentic AI can perform, it seems challenging (albeit not infeasible) to devise datasets to be shared with competitors that are both meaningful and preserve anonymity, as demanded by the DMA. This has already proven to be a challenging task in the enforcement of data sharing of search engines, for which the provision was designed, and where

both the scope and the user-data generated are well defined. Indeed, the datasets to be provided under Article 6(11) DMA have been criticised by competitors for leaving out most of the (relevant) data.

3. FRAND Terms

A final obligation for search engines (as well as app stores and social networks) is found in Article 6(12) DMA, which obliges them to apply "fair, reasonable, and non-discriminatory general conditions of access for business users". Such FRAND obligations are well-known from standardsetting procedures, but they remain a controversial and hard-to-define concept. The obligation of "non-discriminatory" conditions is similar to the ban on self-preferencing and thus comes with similar difficulties (see above). What is "fair and reasonable" needs to be established in relation to the service in question, i.e., AI agents. The issue of enforceability returns here, as the (background) operation of an AI agent is more shielded from the view of users and authorities than that of a traditional search engine.

The implementation of the above, search-engine-specific obligations were already not always as easy to apply to traditional search. Google's implementation of the self-preferencing prohibition of Article 6(5), which is also in Article 6(12), already led to a non-compliance investigation – even after experience with a similar non-discrimination remedy under competition law. Gatekeepers in scope of Article 6(12) generally reported they already complied with its FRAND obligation, which the Commission has not yet questioned. As described above, the sharing of search data has also proven challenging. The implementation and enforcement of these obligations only become more complex for AI agents.

D. Obligations Independent of CPS Category

Finally, the DMA contains a number of obligations that apply to all CPSs, independent of their category. These obligations can be grouped, roughly, in those putting restrictions on data use, those mandating data portability, and those relating to contractual obligations and restrictions.

1. Restrictions on Data Combination and Use

A series of obligations under Article 5(2) DMA are meant to discourage harvesting and crossusing of personal data unless the user has explicitly consented. These provisions, originating from the Bundeskartellamt v. Facebook case, are also rooted in the General Data Protection Regulation (GDPR), and could arguably also be enforced by the responsible data protection authority. The DMA has included these provisions because many of the gatekeepers control entire ecosystems of services, which presents them with a competitive advantage over smaller rivals in collecting, combining and cross-using personal data across services.

Generally, restricting the collection and use of data will also limit the economic efficiencies that are associated with learning from data. Everything else being equal, enabling data sharing (which shares those efficiencies with rivals) is therefore to be preferred over putting restrictions on data use. However, given the complexities of data sharing discussed above, data use restrictions may

be second-best. In this vein, these obligations also seem to make sense in the context of AI agents, which are also likely to be deployed within ecosystems of services of the same provider, including OSs, browsers, search engines, marketplaces and so on.

Further, Article 6(2) DMA restricts the use of "non-public data in competition with business users that was gathered through the activity of business user during their use of the relevant core platform services." Even if designed with marketplaces in mind, this obligation may also be reasonable in the context of AI agents. Specifically, as the AI agent interacts with many thirdparty software applications, which make themselves available for such interaction through offering tools and APIs (as in the example of Apple's App Intents), the agent can potentially learn about user preferences and user behaviour – as the user interacts through the agent with the thirdparty applications – that the provider of the agent may use in competition with those third-party app providers. In similar spirit, with the integration of AI into OSs, OS providers are already collecting data about the use of each individual app.

2. Data Portability

Independent of the CPS type, designated CPS must provide "end users and third parties authorised by an end user" as well as business users effective portability of data according to Articles 6(9) and 6(10) DMA. Specifically, Article 6(9) relates to "data provided by the end user or generated through the activity of the end user in the context of the use of the relevant core platform service" and data access has to be provided continuously and in real-time. As we established above (see Section II), access to such end-user data may be very useful for the training and optimisation of AI agents (and assistants). A popular AI agent has the ability to collect more conversational data and generally more interactions with the user, which can be used to improve the AI agents' outcomes through reinforcement learning. In addition, as AI agents have persistent memory and greater ability to adapt to the user, data portability can counteract data-induced switching costs. Thus, the provision in Article 6(9) is very useful to promote market contestability also in the context of agentic AI.

By contrast, the provision in Article 6(10) does not seem to fit the context of AI agents well. The provision demands that "the gatekeeper shall provide business users … aggregated and nonaggregated data, including personal data, that is provided for or generated in the context of the use of the relevant core platform services". This provision was devised in the context of online marketplaces, where business users' interaction with end users is fully mediated by the platform. The difference to the AI agent context is that the agent acts as a surrogate for the user. Agents thus have the potential to disintermediate the platform and act with business users directly. In this case, there would be no information asymmetry introduced by the platform that is to be corrected through data portability. In the case where AI agents interact with business users via the platform or marketplace, then it is again only the platform (and not the agent) that can meaningfully provide such data portability.

In any case, even when limiting our attention to Article 6(9), the details of the implementation of the provision can become quite complex in the context of AI agents. Even in the seemingly simpler

contexts of "traditional" CPSs, such as social networking, search or messaging, gatekeepers have interpreted "real-time and continuous" data portability as allowing users to schedule daily downloads (once every 24 hours) for a limited period of time. Furthermore, there still exists some legal uncertainty as to the scope of data portability, i.e., what data precisely was "provided by … and generated through the activity of the end user". In order to be meaningful for purposes of contestability, a wide interpretation of the scope of portable data would be helpful. In the context of agentic AI, it is evident that, at a minimum, queries to the AI (prompts) are covered by the data portability provision, but reasonably also the responses by the AI are legitimate "data generated in the context of use" that become subject to portability. In this way, users would be able to port whole conversations with the AI. On the one hand, this would facilitate contestability, as conversational datasets would become portable. On the other hand, widening the scope of data portability may undermine innovation incentives as it promotes the possibility of "distillation", i.e., a process where AI assistants/agents are being trained by extracting data from other models, allowing them to catch up to the incumbent, but without undergoing the same efforts. Nevertheless, along with other observers, we suggest that those responses should be included in the portable dataset. Likewise, data observed by the AI, like sensor data or location, should be subject to data portability. This interpretation of the scope of Article 6(9) would also be

consistent with the corresponding data portability provisions in the Data Act concerning connected products, which explicitly includes "all data generated by the use of a connected product", "including data which result indirectly from the user's action, such as data about the connected product's environment or interactions" and "data generated automatically by sensors".

3. Contractual Obligations and Restrictions

A final set of DMA obligations that apply independent of CPS category relate broadly to contractual obligations and restrictions. As we explain below, however, they have sometimes been adopted with a specific scenario in mind and are not always easily translated beyond that scenario.

A number of obligations are focused on intermediation services. Article 5(3) DMA bans socalled "most-favoured-nation clauses", which have been maintained by hotel booking or marketplace platforms. Insofar as an AI agent intermediates (e.g., for its end-user's purchases), its provider cannot (contractually) prevent business users from offering the same products/services at better conditions elsewhere. As part of this intermediation category, some obligations relate to payments and associated access. Article 5(4) DMA prohibits gatekeepers from maintaining anti-steering policies. Hence, a designated AI agent provider engaging in intermediation could not prevent business users from directly communicating and transacting with end-users. Article 5(5) DMA backs up that provision by obliging gatekeepers to let their endusers access and use, through its CPSs, services/content/subscriptions/features acquired outside of the gatekeeper's CPS. In an AI agent context, this provision gives end-users freedom to invoke related services of their choice, e.g., to specifically buy via the marketplace where they have an account/subscription, rather than the one suggested by the AI agent.

Article 5(7) DMA prohibits gatekeepers from pushing certain services on their business and end-users. The covered services are not CPSs themselves but often form part of a gatekeeper's larger ecosystem. In further support of Articles 5(4)–(5), the provision ensures that end-users can use and business users can interoperate with "a payment service, or technical services that support the provision of payment services, such as payment systems for in-app purchases" of their choosing. The same goes for identification systems. Freedom must also be provided with regard to browser engines. This obligation was adopted for OS gatekeepers that mandate the use of their own browser engine, also for competing browsers. In accordance with the provision, an OS gatekeeper with a browser engine could not prescribe its use by the providers of competing AI agents either. The provision becomes difficult to apply when moving beyond the OS and online intermediation (e.g., app stores, marketplaces) CPS category, even if it theoretically applies there. Consider, for example, a single designation (of an AI agent or different service) as web browser: would its provider also be obliged to offer its users a choice of browser engine in that case? Notably, Article 5(7) does not guarantee freedom of choice regarding AI agents (or virtual assistants), which we return to below (Section VI).

Article 5(8) DMA prohibits tying one gatekeeper service (or even a non-designated service that meets the user thresholds) to another gatekeeper service. Hence, a gatekeeper cannot use access to one service ("tying" service) to reinforce another ("tied" service). The scope of the provision is limited, however, by the fact that the tied service must already have gatekeeper status (or meet the user thresholds); it does not apply to a nascent service. Nevertheless, it can prove useful for defensive leveraging scenarios, where a gatekeeper would want to maintain its strong position in a CPS category by tying that service to one that is in an even stronger position. Depending on its strength, an AI agent could serve as tying or tied product. In the former case,

the concern would be one of foreclosure through an AI agent; in the latter, of an AI agent (see above, Section III).

Article 6(6) DMA, finally, prohibits gatekeepers from restricting the ability from end-users to switch between different apps accessed via their CPSs. The scenario targeted here is one where a gatekeeper in control of an OS or middleware CPS makes user switching between apps more difficult, presumably in favour of the service that is pre-installed. Hence, the provision could promote switching from the gatekeeper's to a third-party AI agent – in other words, it prevents foreclosure of an AI agent (see above, Section III.A). If the AI agent itself is designated, however, the provision would imply that the gatekeeper cannot restrict switching between further downstream services (e.g., marketplaces) – in other words, it prevents foreclosure through an AI agent (see above, Section III.B).

7. SHOULD THE DMA BE REVISED IN THE FACE OF AGENTIC AI?

In this section, we evaluate whether the DMA should be revised in the face of agentic AI. This evaluation is an exercise in synthesis: we start from the competitive concerns (Section III),

checking to which extent they are addressed by obligations in the DMA as they apply to AI agents (Section V). Remember that the applicability of obligations depends on designation: for foreclosure of AI agents, they often need not be designated, as it is the source of foreclosure (likely OSs) that matters; for foreclosure through AI agents, however, designation is more important (Section IV). Our structure here follows the two sets of competitive concerns identified earlier: foreclosure of an AI agent (Section VI.A) vs foreclosure through an AI agent (Section VI.B), with the former being a concern in near term, and the latter in the longer term.

A. Foreclosure of an AI Agent

In the context of foreclosure of an AI agent, we will focus on the most likely scenario, where the gatekeeper of an existing OS, designated under the DMA, pre-installs or otherwise integrates an AI agent on the OS and seeks to foreclose a third-party AI agent on that OS. Without a preinstalled AI agent, the OS provider would have no incentive to foreclose a third-party AI agent. We assume that the pre-installed AI agent is considered the gatekeeper's "own" service, which in our understanding would also be the case if the agent is provided through a close industry cooperation. The pre-installed AI agent may or may not be designated under the DMA in its own right. If the AI agent is designated, we assume it is designated as a "virtual assistant". The designation of the AI agent is especially important for data-related obligations that affect contestability. A separate issue, not addressed here, are contracts between a gatekeeper and another player that involve payments for pre-installing or defaulting to a third-party AI agent.

As noted above (Section III), pre-installation of an AI agent on the OS can give competitive advantages to said agent over third-party agents: users may be reluctant to switch to another AI agent if it is not possible to uninstall the pre-installed one, or if it is not possible to set the new AI agent as the default. Depending on the circumstances, pre-installation can be viewed as a form of self-preferencing (one's own software over others) or a form of tying (shipping one's own software with the OS). When the OS on which the AI agent runs is designated, the DMA's obligations regarding uninstallation and default settings (Article 6(3)) as well as switching (Article 6(6)) and tying (Article 5(8)) generally apply, but the extent to which depends on whether pre-installation concerns the gatekeepers own AI agent, and whether this agent is designated as a virtual assistant CPS in its own right.

Another form of self-preferencing may come in the form of giving one's own AI agent more comprehensive access to the OS, i.e., to the hardware or software features of the device, including more capabilities to collect data on the device (e.g., about the users' location or usage of other apps). Subject to strictly necessary and proportionate integrity measures, Article 6(7) DMA

comprehensively addresses these concerns, i.e., the gatekeeper has to provide equal vertical interoperability to other AI agents, as they provide for their own, and irrespective of whether the own agent is designated as a CPS in its own right. Indeed, Article 6(7) would apply irrespective of whether the pre-installed AI agent is deemed an essential part of the OS or not, as the obligations only refers to "features" available to the gatekeeper. However, provided that AI agents require far-reaching and deep access to the OS to fulfil their tasks in a meaningful way, allowing equal

interoperability is likely to become a contentious issue in enforcement. The closest comparable scenario in today's implementation of the DMA is probably that of enabling alternative browser engines on iOS. Enabling alternative AI agents is likely to be yet more complex, also involving even stronger trade-offs for security and integrity due to their ability to react to user input in various forms (e.g., through visual, audio or motion cues), and their ability perform tasks autonomously on behalf of the user. Therefore, the OS-level access necessitated by AI agents would not only involve those needed by an alternative browser engine (which is required by the AI agents in order to browse the web autonomously), but also additional ones.

The DMA also includes obligations to equate the ability of first-party and third-party service providers to harness and learn from data. This includes the provisions aimed at reducing ecosystemic benefits from data combination and cross-use (Articles 5(2) and 6(2)) as well as the prevention of data-induced lock-in effects through data-portability (Article 6(9)–(10)). These all apply in the context of AI agents shipped with an OS; however, the data portability provisions of the DMA only apply to the pre-installed AI agent directly if it is designated as CPS itself (unlike the data portability obligation of Article 20 GDPR, which applies in any case). Moreover, the data sharing provisions in relation to "click and query" data under Article 6(11) DMA would not apply unless the AI agent is designated as a search engine. Thus, representative, anonymised data (across users) on queries and prompts, as well as users' reactions ("clicks") would not be available to alternative AI agents, even if the gatekeeper AI agent is deemed a virtual assistant CPS, and even though we think this provision would generally be useful in this context (see above, Section V.C.2).

Another form of anticompetitive conduct could revolve around leveraging market power from the OS to promote the use of the pre-installed AI agent by business users. For example, the OS provider could require that software developers (i.e., business users of the OS) integrate features in their software so they work seamlessly with the AI agent (e.g., APIs that facilitate execution of tasks in the app by the agent). This would provide an advantage to the AI agent over other agents that do not benefit from such forced integration. Especially providers of popular apps that are themselves seeking to offer AI agents may be reluctant to integrate competing AI agents in their products, as seen in Meta's unwillingness to integrate Apple Intelligence features. Article 5(7) DMA currently does not cover this case, because it is limited to forced integration with identification services, payment services and web browser engines of gatekeepers. A potential improvement of the DMA in this regard could therefore be to extend Article 5(7) to forced integration with "virtual assistants", whether yet designated as CPS or not.

A final type of conduct covered by the DMA is that of self-preferencing in ranking. While Article 6(5) DMA was designed in reference to search engines, we argued above that the obligation is not strictly limited to this context and may well apply also to the virtual assistant CPS category (see Section V.C.1). This may matter in the present context insofar an OS provider, who ships the OS with a pre-installed AI agent but is also operating a CPS search engine or app store may want to demote rival AI agents in the search results. According to a lawsuit filed by X (née Twitter), this is what Apple is doing to Grok in favour of OpenAI's ChatGPT.165

A set of anticompetitive concerns mentioned above (Section III.A) but not addressed here revolve around foreclosure of off-device bottleneck inputs for the development and functioning

of AI agents, such as foundational model, training data and (cloud) computing resources. Indeed, the DMA has little to say about this. AI models cannot be reasonably mapped to any of the existing CPS categories, nor is there an obligation to share training data of popular services by gatekeepers more generally. Even though cloud computing services are a CPS category, and Article 6(12) DMA obliges certain gatekeepers to apply FRAND access terms, this provision is strictly limited to search engines, app stores and social networking services – it does not apply to cloud computing services. The DMA would therefore not prevent a gatekeeper cloud computing provider from discriminating against specific AI agent providers.

- Obligation covers competition concern.

(□) Obligation covers competition concern with qualification, either because of the nature of the concern (see note * below) or because of the scope of the obligation (see number notes below).

□□□Obligation does not cover competition concern.

Notes:

* Insofar as pre-installation is considered a form of self-preferencing and/or tying.

Obligation only applies if gatekeeper also provides own virtual assistant (not for third-party preinstalled assistants): if gatekeeper's virtual assistant is also designated, a choice screen must be offered; otherwise, only default settings must be changeable.

In relation to on-device data (e.g., GPS data; usage data of installed apps, as it is available to the OS and made available to gatekeeper's own virtual assistant).

Only if the pre-installed virtual assistant is also designated as CPS.

In relation to demotion of the third-party AI agent in search results provided by the OS or, if designated, pre-installed AI agent; and – if applicable – in the gatekeeper's designated search engine and online application store.

Unless the AI agent is designated as search engine.

B. Foreclosure through an AI Agent

Foreclosure of other services through an AI agent is a more long-term concern than foreclosure of an AI agent. This concern only manifests when a provider of an AI agent – whether one of the current gatekeepers or a new player – becomes dominant (or, in the words of the DMA, an important gateway) in its own right. At that point, it may be able to (ab)use that position to foreclose the providers of services downstream of the AI agent. It must have an incentive to do so, which may come from the fact that it provides some of those downstream services itself or that it has agreements granting preferential status to third-party services.

Here, more so than in the previous section, the CPS category in which the AI agent is designated with gatekeeper status matters. Recall that four categories provide a potential fit: virtual assistant, web browser, and search engine in the short- to medium-term; OS in the longterm (see above, Section IV.A

A large number of the obligations can be grouped under the "self-preferencing" heading. Leading this group is Article 6(5) DMA's prohibition of favourable treatment in "ranking and related indexing and crawling". While its terminology clearly originates from search engines, the DMA's broad definition means that virtual assistants (and, if need be, even OSs and web browsers) are covered. The provision also mandates applying "transparent, fair and non-discriminatory conditions" to such ranking. However, providing transparency is technically tricky with regard to AI agents, which also means that establishing fairness and non-discrimination become difficult. Moving on, the FRAND obligation of Article 6(12) DMA applies only to AI agents designated as search engines. There is a partial overlap here with Article 6(5), which covers non-discrimination in ranking. Unless an AI agent is designated as search engine, however, it need not also ensure its access terms are fair and reasonable under Article 6(12).

Another form of self-preferencing is covered by Article 6(3), paragraph 2, DMA: If the provider of a designated AI agent uses it to steer users to its other services (e.g., its marketplace) by default, then it must make that default setting changeable. This obligation only applies, however, if the agent is designated as OS, virtual assistant, or web browser – it does not apply to search engines. The obligation can be seen as complementary to Article 6(5), which does cover search engines, preventing them from steering users to specialised search services by default. While the self-preferencing that motivated the respective obligations – app defaults vs web results – is different, the obligations are formulated so as to cover a wider range of favourable treatment. The important difference is in the exact remedy: While self-preferencing in search is prohibited, app defaults only require changeable settings. The question is how many users will actually change them. There is also a relation with Article 5(8) DMA, which prohibits tying between gatekeeper services. An AI agent could function as tying service, to prop up a different service with gatekeeper status (the tied service), in which case we are talking about foreclosure through an AI agent, but the requirement of two gatekeeper services limits the scope of this provision, even if it does not matter in which category they are designated.

Cognisant of users' status quo bias, lawmakers went a step further than changeable defaults by mandating choice screens for certain (constellations of) CPSs in the second part of Article 6(3), paragraph 2, DMA. While choice screens can remedy foreclosure of an AI agent to some extent (see above, Section VI.A), they have little impact on foreclosure through an AI agent. The obligation only kicks in when an AI agent is designated as virtual assistant or web browser, in which case it would need to show a choice screen but only for search engines. As we discussed above (Section IV.A.2), AI agents do tend to integrate (with) search engines, so the obligation may have merit, even if it is difficult to implement due to the depth on integration with a search engine (and in particular the underlying index).

A final form of self-preferencing, via interoperability, is targeted by Article 6(7) DMA, and may also be relevant for foreclosure through an AI agent. If an AI agent is designated as either OS or

virtual assistant, it needs to provide equal vertical interoperability to other hardware and software providers. A possible scenario here may be that in which the gatekeeper provider of an AI agent also provides a connected device for/with the agent, in which case it would have to grant competing hardware providers the same level of interoperability. As AI agents become system(like) software, interoperability with the agent's software features likely also becomes increasingly important. Note, however, the second part of Article 6(7), which is most important for such access, applies only to OS gatekeepers.

A number of other obligations are more sui generis and have a (theoretically) broad scope. First, the obligation not to use data from business users to compete with them, found in Article 6(2) DMA, would also apply to agents designated in any category, even if it was adopted with online intermediation (marketplaces) in mind. Second, there are the obligations concerning the choices available to business and end-users for communicating and transacting with each other – also outside of the gatekeeper platform – in Article 5(4) and 5(7) DMA. Again, these are generally applicable but kick in primarily when AI agents intermediate, e.g., complete purchases, in which case users should be able to choose how to pay for them.

In conclusion, when it comes to foreclosure through an AI agent, the in our view most fitting qualification as virtual assistant also leads to the best fitting coverage in terms of obligations.

When designing obligations, lawmakers seem to have viewed virtual assistants as (potential)

middleware, i.e., a foundation for other apps or services, which is (coincidentally) in line with the expected development of AI agents. Hence, most of the obligations relating to OSs – and to some extent browsers – also apply to virtual assistants (e.g., changeable defaults/choice screens, equal vertical interoperability). AI agents also exhibit characteristics of search engines (and are in any case integrated with them), but the most important search-related obligations (in particular the prohibition of self-preferencing in ranking) also applies to virtual assistants. A final set of obligations applies independently of the CPS category. Hence, if AI agents are indeed designated as virtual assistants, the DMA does not immediately show gaps when it comes to expected concerns around foreclosure through AI agents. However, as such foreclosure scenarios are further in the future than those regarding foreclosure of AI agents, new gaps may appear over time. By contrast, other CPS qualification that are possible (but not preferred) in the short-term, i.e., the web browser or search engine categories, come with a less comprehensive set of obligations. Interoperability (e.g., of connected devices) would be a problem in the browser designation scenario, while the obligation of changeable defaults does not apply to search engines – to give just two examples.

8. CONCLUSIONS AND RECOMMENDATIONS

Is the DMA future-proof? Its origins and regulatory design do not lead one to expect a positive answer to this question. First, the DMA's origins are found in earlier competition law cases, which means they may be targeted at the issues of yesterday rather than of tomorrow. Second, the DMA

is prescriptive: its scope is determined by a closed list of neatly defined CPSs, and the ensuing obligations are spelled out in similar detail. This is explained by the DMA's origins: given the difficulty of applying open-ended competition law provisions to digital markets, lawmakers took the opposite approach when designing the DMA. But a prescriptive approach – as opposed to, say, a more principles-based one – runs the risk of quickly running into classification problems when technology evolves (does new technology X fit in legal category Y?). Lawmakers did write review provisions with update mechanisms into the DMA, but that makes the regulation adaptive rather than future-proof.

Nevertheless, our case study shows that the DMA proves to be surprisingly future-proof in face of perhaps the most consequential emerging technology of our time – and one that was entirely inexistant during the DMA's drafting: AI agents. The DMA's continued relevance in this new context stems from two facts: one more generally applicable, the other more confined to the subject of our case study. First, the source of foreclosure of AI agents is likely to be found in CPSs that are already designated, particularly (mobile) OSs. This dynamic – where new services rely on established platforms to grow – is more generally applicable, and the DMA is designed to address it. Second, lawmakers added a "virtual assistant" category to the DMA's list of CPSs in trilogue, and its definition captures AI agents. Other CPS categories, such as browsers or search engines, may – due to their broad definitions – also accommodate AI agents, but the fit is not as neat. However, the late addition of virtual assistants – which AI agents can be seen as an evolution of – may be due more to happenstance than prescience. As such, it does not necessarily futureproof the DMA for other emerging technologies, which may instead require adaptation. The DMA's obligations mostly remain effective to address foreclosure of and through AI agents, at least in case of the aforementioned designations.

Notwithstanding the relative suitability of the DMA to tackle potential issues regarding AI agents, some adaptation could help secure its future-proofness. The amendments relate to the

DMA's scope and obligations. Regarding the DMA's scope, we recommend mitigating potential ambiguities in the definition of "virtual assistant". Doing so raises two questions: one of substance (what to change); the other of institutions (how to change). On the substance, we discussed how the definition of "virtual assistant", which is built around the processing of prompts ("demands, tasks or questions"), captures AI agents in the near to mid-term. As AI agents attain ever more capabilities, autonomous action is likely to become more prevalent: the agents may anticipate user "needs" rather than reacting to prompts, which could thus be included in the definition.

As to how to mitigate these ambiguities, we discussed two routes: interpretation (via application) and legislative change. Each route has its benefits in terms of effectiveness, legal (and related investment) certainty, and political legitimacy. A good way to juggle these policy considerations is to view the routes as complements rather than substitutes. In the near to midterm, the designation of AI agents in the virtual assistant CPS category is effective without significantly imperilling legal certainty; in the mid to long-term, legislative change provides a more legally certain way to capture the future iterations of AI agents.

While we favour the designation of agentic AI systems as "virtual assistants" under the DMA in the mid-term, we have also highlighted how the agenticness of AI systems evolves over time. This means that it may be useful to acknowledge this evolution by designating AI systems in different CPS categories over time. Current instantiations of AI assistants show resemblance to search engines, and may therefore considered to be designated as such in the near term. In the long term, AI is likely to become more OS-like, at which point a designation as OS might be more appropriate. As such, our discussion of the various designation possibilities may also prove useful in considering which CPS category may host AI assistants that are not yet on the agentic side of the spectrum.

Regarding the designation of gatekeepers, the existing quantitative thresholds appear sufficient. Companies such as OpenAI, with their early-stage agents, already meet the end-user thresholds. However, complications arise in identifying and counting the business users of virtual assistants. Specifically, developers contacted by AI agents may not actively "use" these agents in the conventional sense. While some developers are taking steps to interoperate with AI agents, and more may do so in the future, others may instead be passively "called upon" by agents. This issue could be resolved through interpretative adjustments aligned with existing practice for search engines. But to increase legal certainty, the Commission should – via delegated act – adapt the DMA's Annex, including its definition of a virtual assistant's business users.

When it comes to obligations, we propose two adjustments so that the DMA adequately covers AI agents. The first adjustment is to extend the click and query data-sharing obligations of Article 6(11) DMA to include virtual assistants. Given the critical role of such data for AI agent contestability, this measure would enhance market openness. The second adjustment is to add virtual assistants alongside browser engines and other services in Article 5(7) DMA. This would help ensure that users retain freedom in their choice of AI agents and, as a result, foster competitive dynamics between such agents.

These proposed amendments aim to counter the foreclosure of AI agents. By contrast, concerns about foreclosure through AI agents do not currently warrant regulatory amendments, assuming AI agents are designated within the virtual assistant CPS category (or in a future AI agent category).

A related concern pertains to the pre-installation of third-party AI agents. While Article 6(3) DMA addresses the pre-installation of a gatekeeper's own services, it does not explicitly cover contractual arrangements with third parties (e.g., partnerships between Samsung and Google, or

Apple and OpenAI), in which there are signs of potential foreclosure. Although this issue lies

somewhat outside the DMA's core logic – which primarily governs internal ecosystem dynamics rather than (contractual) inter-ecosystem relationships – emerging competition among third-party arrangements (e.g., between Motorola and Google, Perplexity and others) suggests that antitrust vigilance is warranted.

Importantly, we have distinguished concerns related to agentic frameworks from those involving foundation models, which we considered only as an input to AI agents. The DMA's provisions do not extend to such models nor their inputs, including training data and compute. As such, the DMA does not ensure non-discriminatory access to off-device computing resources (e.g., cloud services). Given its importance, an idea could be to extend the CPSs covered by the FRAND access obligation of Article 6(12) DMA to include cloud computing services. This would help ensure non-discriminatory access in an increasingly vertically integrated market.

More broadly, structural trends in the AI sector suggest that the DMA alone may not suffice to prevent long-term concentration. Vertical integration across the AI value chain – from chips and OSs to models and agentic frameworks – poses a significant challenge. Due to strong network and scale effects on both the demand and supply sides, increasing market concentration appears likely. This dynamic suggests that rather than seeking for competition in the market, policymakers should strive for competition for the market accompanied with the goal of transient monopolies, in line with Schumpeterian theories of creative destruction.

It is possible that today's gatekeepers – who control key resources like OSs and training data – will also dominate the agentic AI markets. New entrants such as OpenAI could themselves move upstream into hardware and OSs, while incumbent players like NVIDIA may integrate further downstream. However, strategic partnerships of existing gatekeepers with key players in the agentic AI value chain make it likely that the gatekeepers of today are also the gatekeepers of tomorrow. This point relates to the general effectiveness of the DMA; it does not take away from the fact that the DMA is a surprisingly good fit for the advent of agentic AI.

Finally, we acknowledge that applying the DMA to an emerging technology such as agentic AI, whose characteristics are still evolving, involves several trade-offs that must be carefully weighed. Regulating a new technology too early is likely to influence its design and evolution, for better or worse, and the consequences of such intervention can rarely be fully anticipated. In the worst case, regulation may deter innovation and constrain nascent business models in a way that ultimately reduce welfare. By contrast, late intervention offers greater informational clarity: anti-competitive conduct, the appropriate CPS designation, and welfare impacts are more observable once "winners" have emerged. Yet at this stage, strong network effects and user lockin may render remedial measures less effective or even futile. After all, this realization was the very reason for complementing competition law with the DMA's ex ante regulatory framework. While our vision of agentic AI has yet to fully materialize, our case study should at least provide some confidence that the DMA's framework is capable of navigating this trade-off.

References

1. öhsl, J.-F. (2025). Future proofing the DMA for agentic AI: Lessons from the AI Act. World Competition, *48*(3), 315–336.
2. Bostoen, F., & Krämer, J. (2026). How future-proof is the DMA? A case study of AI agents. Journal of Competition Law & Economics, nhag005. Advance online publication. https://doi.org/10.1093/joclec/nhag005

3. Bostoen, F., & Krämer, J. (2025). Is the DMA ready for agentic AI? Centre on Regulation in Europe (CERRE).
4. Colangelo, G. (2026). The DMA meets the new intermediaries: AI agents and the future of gatekeeper regulation (White Paper). International Center for Law & Economics.
5. International Center for Law & Economics. (2026). The DMA's AI dilemma: Too soon, too late, or both? ICLE.
6. DMA High-Level Group. (2025). Joint paper on artificial intelligence. European Commission. https://digital-markets-act.ec.europa.eu
7. Cleary Gottlieb Steen & Hamilton LLP. (2025). Intersection of Digital Markets Act and AI regulation (EU). AI Law - International Review of Artificial Intelligence Law.
8. European Commission. (2025). DMA first review: Consultation on emerging technologies. European Commission.
9. Marinai, S. (2025). Intelligenza artificiale tra intese e collusioni algoritmiche nel diritto dell'Unione europea: questioni interpretative alla luce del DMA, DSA e AI Act. Eurojus.
10. Heinze, C. (2026). Agentic AI in trademark and unfair competition law. Intellectual Property Law Discussion Group, University of Oxford.
11. Rathi, N. (2026, June 24). Re-thinking regulation for the age of AI [Speech]. techUK's Agents of Change: Generative and Agentic AI in Financial Services 2026. Financial Conduct Authority.
12. Goodwin Procter LLP. (2025). Antitrust and AI in the digital age: Antitrust authorities rethink coordination. Insights.
13. European Parliament & Council of the European Union. (2022). Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector (Digital Markets Act). OJ L 265, 12 October 2022, 1–66.
14. European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). OJ L, 2024/1689.
15. European Commission. (2025). Report from the Commission to the Council and the European Parliament: Annual report on Regulation (EU) 2022/1925 (Digital Markets Act) (COM(2025) 166 final).
16. Bostoen, F. (2023). Understanding the Digital Markets Act. The Antitrust Bulletin, *68*(2), 264–290.
17. Franck, J.-U., & Peitz, M. (2024). The Digital Markets Act and the Whack-A-Mole challenge. Common Market Law Review, *61*(2), 299–334.
18. Ibáñez Colomo, P. (2021). The Draft Digital Markets Act: A legal and institutional analysis. Journal of Competition Law & Practice, *12*(7), 573–590.
19. Cremer, J., Dinielli, D., Heidhues, P., et al. (2023). Enforcing the Digital Markets Act: Institutional choices, compliance, and antitrust. Journal of Antitrust Enforcement, *11*(3), 315–350.
20. Tzanaki, A., & Nowag, J. (2024). The institutional framework of the DMA: From hybrid to mature? SSRN Electronic Journal.